

CLAIMS

What is claimed is:

1. A method of mirroring data stored in a source storage system, the method comprising:
 - receiving a plurality of requests at the source storage system;
 - saving modified data in the source storage system based on the requests; and
 - during a synchronization phase, synchronizing data stored in a destination storage system with the data stored in the source storage system, including mirroring at least a portion of the modified data in the destination storage system without requiring said portion of the modified data to be sent from the source storage system to the destination storage system during the synchronization phase.
2. A method as recited in claim 1, wherein said portion of modified data comprises a plurality of blocks of data, and wherein said mirroring at least a portion of the modified data in the destination storage system comprises sending a reference from the source storage system to the destination storage system for each block of said portion of modified data, each said reference for use by the destination storage system to locate the corresponding block in storage controlled by the destination storage system.
3. A method as recited in claim 2, wherein each said reference comprises a transfer ID indicating a data transfer in which the corresponding block was previously sent from the source storage system to the destination storage system.

4. A method as recited in claim 3, wherein each said reference comprises an indication of a location at which the corresponding block was located within the data transfer.

5. A method as recited in claim 3, wherein said mirroring at least a portion of the modified data in the destination storage system comprises storing in the source storage subsystem an association between the transfer IDs and blocks wholly modified by the requests.

6. A method as recited in claim 5, wherein said mirroring at least a portion of the modified data in the destination storage system comprises storing in the destination storage subsystem an association between the transfer IDs and a plurality of offsets, the offsets indicating locations in local storage of the destination storage system at which corresponding blocks of data are stored.

7. A method as recited in claim 1, wherein said portion of modified data consists of blocks wholly modified as a result of the requests.

8. A method as recited in claim 1, further comprising:

creating a log entry in the source storage system for each of the write requests;

and

transmitting each log entry from the source storage system to the destination storage system prior to the synchronization phase, wherein said mirroring at least a portion of the modified data in the destination storage system comprises using data from at least some of the log entries in the destination storage system to mirror said portion of modified data in the destination storage system.

9. A method of mirroring data, the method comprising, in a first storage appliance:

receiving a plurality of requests to write data from a set of client devices, the requests for causing modification of a plurality of blocks of data stored in a first set of non-volatile storage devices coupled to the first storage appliance;

storing modified data in the first set of non-volatile storage devices based on the requests;

initiating a process of synchronizing data in the first set of non-volatile storage devices with data stored in a second set of non-volatile storage devices coupled to a second storage appliance, including

sending each block of a first subset of the plurality of blocks from the first storage appliance to the second storage appliance, to cause the second storage appliance to store the blocks of the first subset in the second set of non-volatile storage devices, and

for each block of a second subset of the plurality of blocks, sending a reference from the first storage appliance to the second storage appliance, instead of sending the corresponding block, each said reference for use by the second storage appliance to locate the corresponding block in local storage of the second storage appliance and to store the corresponding block in the second set of non-volatile storage devices.

10. A method as recited in claim 9, further comprising, prior to initiating the process of synchronizing data:

creating a log entry for each of the requests in the first storage appliance; and

transmitting the log entry for each of the requests to the second storage appliance, by using one or more data transfers, each of the data transfers including one or more of the modified blocks and having a unique transfer ID.

11. A method as recited in claim 10, wherein the second subset of the plurality of blocks comprises blocks which have been wholly modified as a result of the requests.

12. A method as recited in claim 10, wherein said sending a reference from the first storage appliance to the second storage appliance comprises, for each block of the second subset of the plurality of blocks:

 sending a transfer ID and a block number associated with the block to the second storage appliance, the transfer ID identifying a data transfer in which the block was sent to the second storage appliance during said transmitting the log entry, the block number indicating a location of the block within said data transfer.

13. A method of mirroring data, the method comprising, in a first storage server:

 receiving a plurality of requests to write data from a set of client devices, the requests for causing modification of a plurality of blocks of data;

 creating a log entry for each of the requests;

 transmitting the log entry for each of the requests to a second storage server located at a secondary site, using one or more data transfers, each of the data transfers including one or more of the modified blocks and having a unique transfer ID;

 saving modified data in a first set of non-volatile storage devices coupled to the first storage server based on the requests; and

initiating synchronization of data in the first set of non-volatile storage devices with data stored in a second set of non-volatile storage devices coupled to the second storage server, wherein said initiating synchronization includes

for each of the plurality of blocks which has been only partially modified as a result of the requests, sending the partially modified block to the second storage server, and

for each of the plurality of blocks which has been wholly modified as a result of the requests, sending a transfer ID and a block number associated with the wholly modified block to the second storage server instead of the wholly modified block, the transfer ID identifying a data transfer in which the wholly modified block was sent to the second storage server during said transmitting the log entry, the block number indicating a location of the wholly modified block within said data transfer.

14. A method as recited in claim 13, further comprising:

maintaining a transfer ID structure including each said transfer ID;

maintaining a buffer descriptor for each of the blocks;

in response to said transmitting the log entry for each of the requests to a second storage server, storing in the buffer descriptor for each block wholly modified as a result of the requests,

an index to a corresponding transfer ID stored in the transfer ID structure,
and

a block number to indicate a location of the corresponding wholly modified block within a data transfer in which the corresponding wholly modified block was sent to the second storage server during said transmitting the log entry.

15. A method as recited in claim 14, further comprising, in the second storage server:

- receiving the corresponding log entry transmitted from the first storage server for each of the plurality of requests, including receiving the data transfers;
- storing each of the received log entries in local storage of the second storage server, including storing the blocks contained in the data transfers;
- storing each of the transfer IDs of the data transfers in association with a corresponding offset, each offset indicating a location in the local storage of the second storage server at which a block transferred in the corresponding data transfer is stored;
- during said synchronization of data, for each of the plurality of blocks which has been modified as a result of the requests,
 - receiving from the first storage server either a modified block or a transfer ID and block number of a modified block;
 - if a modified block has been received from the first storage server, then storing the modified block in the second set of storage devices; and
 - if a transfer ID and block number of a modified block have been received from the first storage server, then
 - using the received transfer ID to identify the offset associated therewith in the local storage by;
 - using the identified offset to retrieve the modified block from the local storage, and
 - storing the modified block retrieved from the local storage in the second set of storage devices.

16. A system for mirroring data, the system comprising:

a first storage appliance to manage a first set of mass storage devices at a first location, including to receive a plurality of write requests from a set of client devices and to save modified data in the first set of mass storage devices at a consistency point based on the write requests; and

a second storage appliance connected to the first storage appliance via a network at a second location, to manage a second set of mass storage devices, including to mirror at least a portion of the modified data in the second set of mass storage devices at the consistency point without receiving said portion of modified data from the first storage appliance or the first set of mass storage devices at the consistency point.

17. A system for mirroring data, the system comprising:

means for receiving at a source storage system a plurality of write requests;

means for creating a log entry for each of the write requests in the source storage system;

means for transmitting each log entry from the source storage system to a destination storage system;

means for storing each log entry in the destination storage system;

means for saving modified data in the source storage system to non-volatile storage in the source storage system at a consistency point; and

means for using the stored log entries in the destination storage system to mirror at least a portion of the modified data in the destination storage system during a synchronization phase, without sending said portion of modified data from the source storage system to the destination storage system during the synchronization phase, by

sending references to said portion of modified data from the source storage system to the destination storage system.

18. A storage appliance comprising:

- a processor; and

- a memory storing instructions which, when executed by the processor, cause the storage appliance to perform a process that includes:

- receiving a plurality of requests to write data from a set of client devices, the requests for causing modification of a plurality of blocks of data stored in a first set of non-volatile storage devices coupled to the storage appliance;

- storing modified data in the first set of non-volatile storage devices based on the requests;

- initiating a process of mirroring data in the first set of non-volatile storage devices onto a second set of non-volatile storage devices coupled to a remote storage appliance, including

- sending each block of a first subset of the plurality of blocks to the remote storage appliance, to cause the remote storage appliance to store the blocks of the first subset in the second set of non-volatile storage devices, and

- for each block of a second subset of the plurality of blocks, sending a reference to the remote storage appliance, instead of sending the corresponding block, each said reference for use by the remote storage appliance to locate the corresponding block in local storage of the remote storage appliance and to store the corresponding block in the second set of non-volatile storage devices.